DOCUMENT RESUME

ED 145 837 IR 005 278

AUTHOR Sinnott, Loraine T.
TITLE Generative Computer-Assisted Instruction and Artificial Intelligence. Report No. 5.
INSTITUTION Educational Testing Service, Princeton, N.J.
SPONS AGENCY Advanced Research Projects Agency (DOD), Washington, D.C.
PUB DATE Oct 76
CONTRACT MDA-903-74-C-0290
NOTE 31p.

EDRS PRICE MF-$0.83 HC-$2.06 Plus Postage.
DESCRIPTORS *Artificial Intelligence; *Computer Assisted Instruction; Instructional Technology; *Models; Programing Languages; *State of the Art Reviews

ABSTRACT

This paper reviews the state-of-the-art in generative computer-assisted instruction and artificial intelligence. It divides relevant research into three areas of instructional modeling: models of the subject matter; models of the learner's state of knowledge; and models of teaching strategies. Within these areas, work sponsored by Advanced Research Projects Agency (ARPA) plays a prominent role in theoretical advances. ARPA sponsored research has concentrated both on the development of instructional systems teaching subjects like programming and problem solving, and on research basic to the improvement of adaptive instructional systems in general. The review concludes that as a CAI system becomes more responsive to natural language input, the number of extraneous skills a student must develop in order to interact with the program decreases. Also, providing an author with the opportunity to interact with the computer in natural language lessens the time required to create CAI materials, as well as the constraints imposed on those materials by working in a programming language. (Author)

****
* Documents acquired by ERIC include many informal unpublished *
* materials not available from other sources. ERIC makes every effort *
* to obtain the best copy available. Nevertheless, items of marginal *
* reproducibility are often encountered and this affects the quality *
* of the microfiche and hardcopy reproductions ERIC makes available *
* via the ERIC Document Reproduction Service (EDRS). EDRS is not *
* responsible for the quality of the original document. Reproductions *
* supplied by EDRS are the best that can be made from the original. *
****

ED145837

U S DEPARTMENT OF HEALTH,
EDUCATION & WELFARE
NATIONAL INSTITUTE OF
EDUCATION

THIS DOCUMENT HAS BEEN REPRODUCED EXACTLY AS RECEIVED FROM THE PERSON OR ORGANIZATION ORIGINATING IT. POINTS OF VIEW OR OPINIONS STATED DO NOT NECESSARILY REPRESENT OFFICIAL NATIONAL INSTITUTE OF EDUCATION POSITION OR POLICY

Report No. 5

GENERATIVE COMPUTER-ASSISTED INSTRUCTION AND ARTIFICIAL INTELLIGENCE

Loraine T. Sinnott
Educational Testing Service
Princeton, New Jersey 08540

October, 1976

APPROVED FOR PUBLIC RELEASE
DISTRIBUTION UNLIMITED

Sponsored by

Advanced Research Projects Agency
1400 Wilson Blvd.
Arlington, Virginia 22209

ARPA Order No: 2651
Contract No. MDA 903-74-C-0290

IR005278

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER: 12	2. GOVT ACCESSION NO
3. RECIPIENT'S CATALOG NUMBER	
4. TITLE (and Subtitle): Generative Computer-Assisted Instruction and Artificial Intelligence	5. TYPE OF REPORT & PERIOD COVERED: Report No. 5
	6. PERFORMING ORG. REPORT NUMBER: MDA 903-74-C-0290-12
7. AUTHOR(s): Loraine T. Sinnott	8. CONTRACT OR GRANT NUMBER(s): MDA 903-74-C-0290
9. PERFORMING ORGANIZATION NAME AND ADDRESS: Educational Testing Service, Princeton, New Jersey 08540	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS: 4W10-0002
11. CONTROLLING OFFICE NAME AND ADDRESS: DARPA, 1400 Wilson Blvd., Arlington, Virginia 22209	12. REPORT DATE: October 1976
	13. NUMBER OF PAGES: 29
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office): Same	15. SECURITY CLASS. (of this report): Unclassified
	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

APPROVED FOR PUBLIC RELEASE
DISTRIBUTION UNLIMITED

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Adaptive Instruction
Artificial Intelligence
Computer-Assisted Instruction
Generative Computer-Assisted Instruction
Instructional Models
Models of the Learner
Models of the Subject Matter
Teaching Strategies

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This paper reviews the state-of-the-art in generative computer-assisted instruction and artificial intelligence. It divides relevant research into three areas of instructional modeling: models of the subject matter; models of the learner's state of knowledge; and models of teaching strategies. Within these areas, work sponsored by ARPA plays a prominent role in theoretical advances. ARPA sponsored research has concentrated both on the development of instructional systems teaching subjects like programming and problem solving, and on research basic to the improvement of adaptive

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

instructional systems in general. The review concludes that as a CAI system becomes more responsive to natural language input, the number of extraneous skills a student must develop in order to interact with the program decreases. Also, providing an author with the opportunity to interact with the computer in natural language lessens the time required to create CAI materials as well as the constraints imposed on those materials by working in a programming language.

The views and conclusions contained in this document are those of the author and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Advanced Research Projects Agency or the United States Government.

CONTENTS

	Page
INTRODUCTION	1
PERSPECTIVE	2
STATE-OF-THE-ART AND ARPA RESEARCH	4
Modeling the Subject Matter	4
Modeling the Learner's State of Knowledge	13
Models Used for Diagnostics	13
Models Used to Record Student Knowledge	15
Models of the Process of Acquisition	17
Modeling Teaching Strategies	18
CONCLUSIONS	19
TRENDS	20
REFERENCES	22

Generative Computer Assisted Instruction and Artificial Intelligence

INTRODUCTION

As has often been noted (e.g., Carbonell, 1970a; Uttal, Pasich, Rogers, & Hieronymous, 1969), traditional frame-oriented computer-assisted instruction does not make full use of the capabilities of the computer. In frame-oriented CAI students are presented prestored text and text-related questions. Although complex branching strategies may be employed, frame-oriented CAI is basically a mechanized version of a programmed text. As such it makes maximal use of the computer's capacity for storing large amounts of data, but minimal use of the computer's logical capabilities.

Certain modes of computer-assisted instruction, such as tutorial and drill-and-practice, have been successfully implemented in frame-oriented environments (Rockart & Scott Morton, 1975). But frame-oriented systems cannot effectively handle other forms of CAI applications, such as those supporting problem solving activities or explorations of social or physical system simulations. These forms require a program which can do more than literal manipulation of text in sequences predetermined by the program's author; they require a system which can adapt to specific, unanticipated interactions with the student.

The class of CAI programs termed "generative" encompasses systems which can interact responsively in less predictable modes of CAI, like problem solving or computer simulations. Through programmed procedures, generative systems bring together elements of data bases to construct, for example, answers to questions posed by the student, segments of instructional text, questions directed to the student, or answers to computer constructed questions.

A trend in the development of generative CAI has been to develop systems which act much like intelligent human tutors. Since a goal of artifical intelligence research is "to construct programs which exhibit what we call intelligent behavior" (Feigenbaum & Feldman, 1963), AI is the breeding ground for techniques applied in generative systems which aid in the

achievement of this tutor-like environment. Implementation of techniques from AI has produced a new level of sophistication in CAI applications, providing environments only partially described by the term "generative." A more descriptive designation for these systems, stemming from their emphasis on tutor modeling, is "responsive" computer-assisted instruction.

PERSPECTIVE

In 1969, Uttal et al. described generative programs as those which use "algorithms to generate problems, answers, diagnostics, and remedial materials." At the time, they felt that not all subject domains were amenable to generative development, just those conducive to algorithmic manipulation, such as "mathematical course materials, other mathematically oriented subjects like the physical sciences, and some special subject matters such as chemistry or logic in which there is a formal system describing the relations between and among parts." In particular, they excluded "verbally oriented subject matter as possible items for a generative curriculum."

A year later, two systems emerged which indicated that verbally oriented subject matter could be taught in a generative environment (Carbonell, 1970a; Wexler, 1970). Both were tutors of geography. Though applicable to other subject areas, both were oriented toward the tutoring of factual knowledge.

The systems used structured information networks to store key concepts in their subject domain. But the locus of the meaning of those concepts differed between the programs. Meaning in Wexler's system resided, for the most part, in a prestored set of incomplete sentences. Concepts were retrieved from the network to appropriately complete the statements, which were then transmitted to the student.

The emphasis of meaning in Carbonell's system resided in the network itself. Following the work of Quillian (1968) in natural language comprehension, Carbonell constructed a detailed network of interrelated concepts stored in attribute-value (e.g., Capital-Santiago) format. Facts about a geographical concept (e.g., Argentina) could be inferred from the network by procedures

which interpreted a list of properties associated with the concept, each property consisting of an attribute and its value or set of values. The values in turn had their own associated property lists, allowing for multiple embedding of information. No pieces of text or questions were part of the network. Procedures, almost completely independent of the subject matter to which they applied, were capable of searching the network to construct the output the system presented to the student.

Although Wexler provided a generative tutor in a subject considerably less formal than those Uttal and his associates thought possible, his system was not a theoretical departure from the notion of algorithmic generation held by these investigators. On the other hand, Carbonell's signified a new direction in the evolution of generative programs. By endowing SCHOLAR, as the system was called, with a representation of its subject matter in the form of a network and procedures which understood it, Carbonell took the step toward providing CAI programs with independent problem solving and question answering expertise. Carbonell borrowed techniques from artificial intelligence, specifically from the area of natural language comprehension, to give SCHOLAR some 'understanding' of geography.

Since SCHOLAR's development, applications of AI in CAI have continued, with exemplary programs including a logic teaching system (Goldberg, 1973), a system for teaching integration (Kimball, 1973), a tutor of a text-editing system (Grignetti, Gould, Hausmann, Bell, Harris, & Passafiume, 1974), a laboratory in electronic troubleshooting (Brown, Burton, & Bell, 1974a), a tutor of a machine language (Koffman & Blount, 1974, 1975), and a tutor of BASIC (Barr, Beard, & Atkinson, 1975a). The following statement by Brown, Bell, and Zdybel (1973) reflects the underlying philosophy of this approach:

> It seems almost axiomatic that the more a CAI system 'understands' about its subject area, the more effective a teaching system it will be. Even more important is the fact that systems which possess such understanding can provide a qualitatively new level of interaction with the student.

STATE-OF-THE-ART AND ARPA RESEARCH

ARPA has supported research in the development of responsive CAI environments. The research has produced both practical and theoretical accomplishments. On the practical side can be listed the development of tutors which have the potential to effectively and efficiently train military personnel. On the theoretical side can be listed advances in the modeling of subject matter, the modeling of the learner's state of knowledge, and the modeling of teaching strategies. The following ARPA supported projects will be discussed as they relate to these accomplishments:

(1) The work of Brown and associates at Bolt, Beranek, and Newman, Inc., on instructional environments for problem solving and gaming.

(2) The work of Collins and associates, also at BBN, on research with SCHOLAR-like systems.

(3) The work of Atkinson and associates at Stanford University on an instructional system for teaching the BASIC programming language.

and (4) The work of Norman and associates at the University of California at Irvine on human information processing.

Modeling the Subject Matter

The capability of an automated tutor to respond to the needs of its student is greatly enhanced if the tutor can interpret and pertinently reply to unanticipated student input. If the tutor is capable of this, the student can, for example, ask for clarification of some previously discussed point or fill a critical gap in his knowledge by querying the tutor. Furthermore, the range of student answers to questions the computer poses need not match a finite set of prestored statements; the computer can provide pertinent responses to many unforeseen answers.

Programs which can flexibly and intelligently handle unanticipated input generally require a representation of the content of their subject area to which they can turn to interpret the meaning of the input and to construct

a relevant response (Nilsson, 1974). This requirement has inspired work on modeling representations of knowledge for use in automated tutors. Subject matter modeling has typically been directed at modeling specific subject areas, but more general representations have been attempted (e.g., Kingsley & Stelzer, 1974).

SOPHIE, a tutor developed by Brown and coworkers, can meaningfully react to unanticipated ideas concerning what is wrong with a faulted piece of electronic equipment (Brown & Burton, 1975). Basic to SOPHIE's ability to do this are several representations of the subject area, each of which supports different aspects of SOPHIE's interactions with students. Included in the representations called upon by SOPHIE are (a) a semantically-centered parser of student input, (b) a network which encodes time-invariant factual knowledge about the instrument's circuit, (c) a model which can simulate the behavior of the instrument, and (d) an assortment of heuristic procedures, each of which specializes in inferring specific types of information from the results of the circuit simulation. Before discussing these representations in more detail, a description of the purpose and appearance of SOPHIE's tutorial interaction will be given.

SOPHIE's goal is to teach the qualitative understanding of electronic circuits necessary for troubleshooting. To this end SOPHIE interacts with the student while the student debugs a malfunctioning piece of electronic equipment. Its success at supporting this interaction has led to SOPHIE's evaluation as the most impressive artifical intelligence-based system yet developed (Bunderson & Faust, 1976).

The tutorial begins with the presentation of a schematic of the component's circuit and a partial specification of the equipment's symptoms for some particular but unknown fault. The student is then permitted to examine the behavior of the instrument by requesting certain measurements of its output under various control settings or requesting tests at internal points of the circuit. The student can ask SOPHIE specific questions about the implications of a measurement. He can ask SOPHIE advise about what to consider next given what has thus far been discovered about the circuit. SOPHIE can evaluate a student's hypothesis about what is wrong with the

circuit, taking into consideration all the information the student should have been able to derive from his previous measurements. SOPHIE can also judge the merits of any particular measurement request in view of the prior sequence of measurements the student has made.

A portion of a SOPHIE-student interaction taken from Brown, Burton, Bell and Bobrow (1974) is presented in Figure 1. Although the dialogue is difficult to fully appreciate without some prior knowledge of electronic circuits, it should still provide some idea of how flexible and seemingly intelligent SOPHIE is.

Figure 1

Segment of a SOPHIE-student interaction taken from a longer protocol presented in Brown, Burton, Bell, and Bobrow (1974). Statements entered by the student are underlined.

WHAT IS THE BASE EMITTER VOLTAGE OF THE VOLTAGE LIMITING TRANSISTOR?

THE BASE EMITTER VOLTAGE OF Q5 IS -.Ø37 MILLIVOLTS.

WHAT ABOUT FOR THE CURRENT LIMITING TRANSISTOR?

THE BASE EMITTER VOLTAGE OF Q6 IS .349 VOLTS.

WHAT SHOULD IT BE?

IN A WORKING CIRCUIT THE BASE EMITTER VOLTAGE OF Q6 IS .593 VOLTS.

REPLACE Q6.

I AM GOING TO ASK YOU SOME QUESTIONS ABOUT HOW Q6 IS FAULTED.
ARE ANY JUNCTIONS SHORTED?

I DON'T KNOW.

ARE ANY TERMINALS OPEN?

YES.

WHICH TERMINALS ARE OPEN?

THE BASE.

Q6 IS NOT FAULTED IN THAT MANNER. Q6 NOT REPLACED.

IS IT POSSIBLE THAT Q6 HAS AN OPEN BASE?

IS THE COMPLETE FAULT MODE EB/EC/BC = OP/OK/OP WHAT YOU INTENDED?

SURE.

WHEN THE SETTINGS WERE: CC=1.Ø, CR=HIGH, LOAD=2Ø, VC=1.Ø, VR=HIGH
THAT FAULT CONFLICTS WITH THE FOLLOWING MEASUREMENT. YOU OBSERVED
THAT THE OUTPUT VOLTAGE WAS 11.7 VOLTS. IF Q6 HAD FAULT MODE
EB/EC/BC = OP/OK/OP IT WOULD BE 23.9 VOLTS. IN A WORKING CIRCUIT
IT IS 19.9 VOLTS. ALSO, THAT FAULT CONFLICTS WITH THE FOLLOWING
MEASUREMENT. YOU OBSERVED THAT THE BASE EMITTER VOLTAGE OF Q6 WAS .349
VOLTS. IF Q6 HAD FAULT MODE EB/EC/BC = OP/OK/OP IT WOULD BE .712 VOLTS.
IN A WORKING CIRCUIT IT IS .593 VOLTS. THAT FAULT DOES NOT EXPLAIN ANY OF
THE SYMPTOMS YOU OBSERVED BUT IT IS CONSISTENT WITH THE VALUES FOR THE BASE
EMITTER VOLTAGE OF Q5 AND THE VOLTAGE AT N1.

In SOPHIE, student input is converted to executable LISP code by a semantically based parser (Brown, Burton, Bell & Bobrow, 1974; Burton, 1974); evaluation of the code provides the data which comprises SOPHIE's response to the input. A semantically based parser is to be distinguished from a syntactically based parser (like, for example, an augmented transition network parser (Woods, 1970)). In syntactic parsers, substrings of input are associated with syntactic categories such as "noun", "verb", "noun phrase" or "adjective". In semantic parsers substrings are related to conceptual categories such as (in the case of SOPHIE) "a request for information", "a command to change a control setting", "a measurement" or "a location".

In general a semantic parser could lead to a phenomenal proliferation of categories to be captured by the grammar. But in the area of electronic troubleshooting the number of concepts is limited enough and the constraints that interrelate them are well defined enough to make the use of such a parser feasible. The constraints built into the grammar between the various conceptual entities which SOPHIE must comprehend represents one dimension of SOPHIE's multidimensional "knowledge" of its subject matter.

SOPHIE derives its factual knowledge about the circuit from a network representation of this information. The network stores only time-invariant information such as the specifications of circuit components or definitions of circuit terminology. The structure of the network is in the tradition of Quillian (1969). SOPHIE calls upon the network when it answers factual questions about the circuit posed by the student. The network is also called upon by SOPHIE's heuristic procedures and natural language processor when factual information is required for their operation.

But SOPHIE must also deal with time-variant information about the circuit since the circuit's state is constantly undergoing modifications during the course of troubleshooting. It would be impossible to provide SOPHIE with an explicit representation of all potential circuit states since there are an indefinite number of them. Instead SOPHIE has access to an implicit representation of all potential states: a model which simulates the real circuit's behavior. When SOPHIE requires information about the circuit in some specified condition, SOPHIE activates the simulation to produce a circuit representation in the required state. SOPHIE then has at its disposal a description of the modified circuit from which it may gather the information it needs. Grignetti et al., (1974) draw the following analogy when observing how natural it is to use a simulation model in a CAI application: "A person's data base is not only memory, and his 'retrieval routines' are not solely introspective: he uses the world as a data base and his senses to retrieve information from it. I don't need to have in my head a representation of what is behind my chair; if I need to know, I can just turn around, look, and see!"

According to Brown, Burton, and Bell (1974), "the main seat of intelligence in SOPHIE resides in its ability to draw conclusions and make inferences from setting up, running, and examining the simulation model of its problem domain." The availability of the circuit simulation and heuristic procedures which can interpret it is responsible for the success with which SOPHIE handles the following communications: answers to questions about the state of the circuit after some modification of the circuit, evaluations of student hypotheses concerning the equipment's fault, and suggestions of hypotheses that the student might consider.

More generally, for subject domains amenable to such modeling, the use of a simulation can equip a tutor with the means to make powerful inferences. For domains of knowledge within which simulation models can be constructed, "we can create an artifically intelligent 'expert' system which can patiently provide the student with a logically deep sounding board for his own ideas (Brown, Burton, & Bell, 1974)." Subject areas within which simulation models can be constructed include (Brown, Burton, & Bell, 1974):

(1) programming, wherein one can simulate the execution of a program to understand how it works,

(2) areas of medical science within which physiological or pharmacological processes can be modeled,

and (3) areas of engineering within which physical systems can be modeled.

Work by Collins and associates suggests that powerful inferential mechanisms may also be available in subject areas whose content is not adaptable to partial or total representation by simulation models. These investigators have developed procedures which search a semantic network a la Carbonell in order to derive implicit facts consistent with or suggested by the explicit information which the network houses. Collins and associates are working with SCHOLAR-like tutors and have, in fact, continued work on Carbonell's geography tutor, now called GEO-SCHOLAR. They have also developed SCHOLAR-like tutors of the ARPA network (Grignetti & Warnock, 1973) and a text-editing system called NLS (Grignetti et al., 1974). Before discussing the inferential procedures that Collins and coworkers have modeled, a description of the purpose and appearance of SCHOLAR's tutorial interaction will be given. GEO-SCHOLAR is chosen to exemplify the family of SCHOLAR-like tutors and will be the focus of our discussion.

GEO-SCHOLAR can handle unanticipated student dialogue in a subset of English which is constrained, for the most part, by SCHOLAR's limited syntactic capabilities. The content of the dialogue is dependent on student responses, questions, and requests. SCHOLAR uses its semantic network to generate the text it presents, the questions it asks, and the responses it constructs to both its own questions and those posed by the student. SCHOLAR can call upon programmed procedures to make several types of calculations and inferences.

Figure 2 presents a sample protocol of an interaction with GEO-SCHOLAR. It is taken from Collins and Warnock (1974).

Figure 2

A GEO-SCHOLAR-student interaction. The protocol was originally presented in Collins and Warnock (1974). Statements entered by the student are underlined.

WHAT IS THE LANGUAGE OF ARGENTINA?

PROBABLY SPANISH.

VERY GOOD.
APPROXIMATELY WHAT IS THE AREA OF ARGENTINA?

TELL ME SOMETHING ABOUT PERU.

PERU IS A COUNTRY. IT IS LOCATED IN SOUTH AMERICA. THE CAPITAL OF PERU IS LIMA.

THE AREA OF ARGENTINA IS 1,100,000 SQUARE MILES.

APPROXIMATELY TRUE, BUT A MORE EXACT VALUE IS 1,200,000.
THE EASTERN BORDERING COUNTRY OF PERU IS _ _ _ _ _.

CATHOLICISM.

WRONG. CATHOLICISM IS NOT A COUNTRY BUT A RELIGION. THE CORRECT ANSWER IS:
BRAZIL. WHERE IS PERU?

WHAT IS THE MOST IMPORTANT DIFFERENCE BETWEEN PERU AND PARAGUAY?

THE PRIMARY DIFFERENCE IS: THE POPULATION OF PERU IS LARGER BY A FACTOR OF 7.8.

Collins and associates have attempted to extend SCHOLAR's information base beyond that which is literally stored in the semantic network in order to broaden both the range of knowledge about geography which SCHOLAR can talk about and respond to, and the types of knowledge SCHOLAR can effectively teach. They began this effort by classifying the inferential strategies used by people (Carbonell & Collins, 1973). Several of the strategies identified have been implemented as procedures used by SCHOLAR (Collins & Warnock, 1974; Collins, Warnock, Aiello & Miller, 1975).

To illustrate the kinds of inferences SCHOLAR can make, several examples will be described. The examples are presented in a positive, fail safe manner but the reader should keep in mind that the success of an inference depends on the presence and appropriateness of facts in the network upon which the inference rests.

(1) When a property is not on the property list of a concept, SCHOLAR can still determine whether the concept has the property. For example, to answer the question "Does the Llanos have a rainy season?" SCHOLAR first searches through the information directly linked with the Llanos. With the present data base, it finds nothing about a rainy season. It does find, though, that the Llanos is a savanah. Searching through the property list of "savanah", SCHOLAR finds that a savanah has a rainy season. SCHOLAR then infers that the Llanos must also.

(2) SCHOLAR can determine whether a question is true or false. For example, to answer the question "Is Bolivia a capital?" SCHOLAR observes that Bolivia is a country and that a capital is a city. It also has stored the fact that a city and a country are different kinds of places. Hence SCHOLAR answers "no".

(3) SCHOLAR can determine the agricultural products or climate of a geographical region. For example, to answer the question "Is the climate of Buenos Aires subtropical?" SCHOLAR first checks the property list of Buenos Aires but finds no mention of it being subtropical. SCHOLAR then attempts to locate a city which is subtropical. It finds Caracas. SCHOLAR then checks whether the latitude and longitude of Caracas are similar to those of Buenos Aires. They are not. SCHOLAR infers that it is unlikely that Buenos Aires is subtropical.

The strategy SCHOLAR employs to answer questions about products and climate depends on SCHOLAR's endowment of knowledge about the determinants of products (climate, rainfall, and soil fertility) and climate (longitude and latitude). It also depends on the general, and sometimes misleading, rule that if certain criteria of similarity are met between two situations then a result that pertains to the first pertains to the second.

Using inferences like those illustrated above, SCHOLAR can now talk about information that is not explicitly stored in the network. This power has greatly enhanced the system's ability to interact with students both by enlargening the range of student inputs it can intelligently address and by increasing the depth of its responses.

The networks used by SCHOLAR and SOPHIE are instances of a class of structures known as semantic networks. As the name implies, a semantic network attempts to represent the meaning of facts. Typically links associate the facts into a totally connected structure. In the case of SCHOLAR, for example, attributes can be thought of as the links of the network. Other examples of semantic networks can be found in the works of Anderson and Bower (1973), Norman and Rumelhart (1975), and Simmons (1973), while a discussion of fundamental issues is presented by Woods (1975).

The work of Norman and Rumelhart (1975) (and also the work of Anderson and Bower, 1973) on semantic networks is the product of an attempt to model human memory. Because of this, a discussion of the Norman-Rumelhart model will be given in the next section on modeling the learner. However we would like to briefly discuss at this point an extension of their modeling effort into the realm of representing complex subject matter.

Semantic networks as typically designed readily represent factual knowledge, but confront a more difficult problem in attempting to model the higher level conceptualizations which exist in most subject areas. For example, "it is easy to represent that General Grant smoked cigars or that the Amazon River had numerous tributaries, but not easy to represent that the North's plan of action in the Civil War was to cut off the South from external supplies, and that this plan motivated Lincoln to keep a major segment of troops at Cairo, Illinois and St. Louis, and that it eventually led Grant to cut off the Tennessee and Cumberland rivers (Norman, 1976)."

Drawing from work by Rumelhart (1975, 1976) on representing simple stories, Norman and Rumelhart (Norman, 1976) are developing a scheme to represent complex episodes, like the Civil War episode given above. A tutor of history is presently under development by these investigators that will provide a vehicle within which their representation ideas can be implemented and evaluated. Although the resulting model will apply

directly to historical narratives, its development may also provide some understanding of how to structure the complex interrelations which comprise other subject areas.

Modeling the Learner's State of Knowledge

As was suggested in the previous section, if a CAI program has access to information that endows it with some subject matter expertise, decisions about what material to present can be made dynamically as the student progresses through the course. Tutor-tutee dialogue thus has the potential to be keyed to the needs of the student. But a model of its subject domain provides only part of the data the tutor may require to respond adaptively to the student; the tutor may also require information about the student. As Norman (1973) has noted, without a model of the knowledge of its student, a computer based system that supports interactive dialogue has only reached the stage of egocentric intellectual development.

Three kinds of models of the learner which may be implemented within a tutor will be illustrated. They differ primarily in how they function within the tutor. We discuss first models employed by the tutor to diagnose student errors. Such models typically depend on taxonomies of errors the student may make. We discuss next models employed to record what the student has learned about the subject area. Such models typically depend on explicit representations of those parts of the subject matter the student has mastered. And finally, we discuss models employed to determine the manner in which information should be presented. Such models are typically operational statements of how students acquire knowledge.

Models Used For Diagnostics

Assuming a student's answer to a question requires intermediate steps for its derivation, when the answer is wrong the tutor can choose to attend to the answer, its derivation, or both. For example, if the tutor acknowledges that the answer is incorrect and provides the right answer, the tutor has attended to the product--the answer. On the other

hand, attention to the derivation is achieved if the tutor indicates faults in the student's reasoning or begins to question the student about the manner in which he arrived at his answer.

The tutor can attend to the derivational process without concern for the particular process employed by the student. For example, Wexler's (1970) geography tutor responded to a student's incorrect responses by outputting the derivational path the computer would have considered had it been asked the question. But this form of remediation is not tailored to the weaknesses of the student.

If remediation is to be keyed to the student's derivational process, the tutor must turn to procedures which recognize faults in the student's reasoning. The process of recognizing faults is complicated by the fact that it is often unnatural for a student to detail every step he has taken to arrive at his answer. Hence the procedures may have to infer mistakes in reasoning from a, perhaps, sketchy account of the steps the student has taken. Whatever inference strategies are employed though, the process of diagnosing student faults is likely to depend on the presence within the system of a taxomomy of potential errors the student may make (Carbonell, 1970b; Siklossy, 1970).

Brown and coworkers employ taxonomies in both a tutor which provides guidance while a student plays a game (Burton & Brown, 1976) and a tutor which provides guidance while a student solves algebraic equations (Brown, 1975a, b). The game tutor points out weaknesses in the strategy a student employs to play the PLATO game "How the West Was Won",* a game supporting drill and practice in elementary arithmetic. The algebra tutor helps a student acquire algebraic skills by presenting him with problems and providing him with feedback relating to his solutions of the problems.

---

*This game was written by Bonnie Anderson for the PLATO Elementary Mathematics Project.

The algebra tutor bases its diagnostics of problems incorrectly solved on a taxonomy of errors which reflect universal kinds of misunderstandings related to algebraic manipulations. A subset of the taxonomy the tutor uses is given in Brown (1975a). The taxonomy used by the game tutor is not really a taxonomy of errors, but a taxonomy of strategies or basic operations the student may employ to play the game. The diagnostic module calls on this taxonomy of strategies to note what operations the student employs in his moves and what operations the student could employ to improve his game. In both tutors, remediation is based on weaknesses which the diagnostic procedures uncover.

Both tutors rely on records of past performance to more accurately determine the weaknesses in a student's performance. Student histories are used to cope with the difficulty suggested earlier, namely the difficulty in inferring faults based on only a few isolated student inputs. In the game tutor, a student's single move typically does not provide enough information to conclude that he is unaware of some basic strategy that would improve his game. The game tutor calls upon a summary of the student's previous moves to aid in the inference of deficiencies. Similarly, in its attempt to ascertain the improper operations the student has used to solve a problem, the algebra tutor turns to a summary of problem solving techniques presently in the student's repertoire. Included in the summary are both the legitimate, heuristically sound transformations known to the student and the illegal or counter productive transformations he has employed and may employ again.

Models Used To Record Student Knowledge

The summaries used by Brown's tutors to aid in diagnosing student weaknesses are instances of models which record the state of a student's knowledge. In the case of the game tutor the student's state of knowledge is defined by his repertoire of strategies; in the case of the remedial algebra tutor, the student's state of knowledge is defined by his repertoire of algebraic manipulations. Models which reflect a student's present understanding of the subject area can provide data influencing a number of tutor activities, such as decisions concerning what to teach next or the depth of remediation. Such models can be used to structure the tutor's explanations in terms familiar to the student. They can also provide data pertinent

to determining why a student has asked a specific question, or, as we have seen in the section on diagnostics, determining why a student has incorrectly answered a question.

As in research on modeling the subject matter, a variety of structures which model the student's present state of knowledge have been proposed. Indeed, the two modeling efforts are intimately connected. Methods used to model the subject matter may be applicable to modeling the student's state of knowledge. For example, Self (1974) has proposed that some forms of knowledge can be readily represented as programs. While the representation of the subject domain can consist of programs which, when computed, result in the "ideal" answers, those which model the student produce answers which match those the student would supply.

The model of the subject area itself may provide a natural depository for data regarding what the student has learned. Both Wexler (1970) and Carbonell (1970a, b) suggested flagging their networks in order to keep track of areas the student had mastered. Atkinson and associates have also taken this approach in their tutor of BASIC known as BIP (Barr, Beard, & Atkinson, 1975a, 1975b; Beard, Barr, Fletcher, & Atkinson, 1975).

BIP (BASIC Instructional Program) is an interactive problem-solving laboratory that offers tutorial assistance to students solving introductory programming problems. The system is a stand-alone, fully self-contained course in BASIC programming. The emphasis of the course is on solving programming problems. To this end students are presented a series of programming problems which they solve on-line. As the student develops a program he is directed to appropriate sections of a hard-copy manual which explains such things as BASIC statements and programming structures. The student is also encouraged to use numerous on-line student-oriented features such as interactive debugging facilities and help options.

To aid in decisions concerning the sequencing of programming problems, BIP calls upon a Curriculum Information Network (CIN). A CIN is a structured representation of the curriculum which consists, in part, of the problems or tasks the tutor may present to the student, linked to the basic skills required for successful completion of the problems. BIP uses this network to flag the skills its student has mastered. Presently each skill

is tagged with an array of counters which note the student's state of acquisition of the skill. The counter values are determined by the student's ability to complete programming problems which require the skill and the student's own assessment of his mastery.

Models of the Process of Acquisition

Models of how students acquire knowledge are basic to the understanding of how information should be presented. If operational models were available, the tutor would have a powerful means for determining the structure of its tutorial interaction. In the area of paired-associate learning, a variety of mathematical models exist which describe probablistically how information is acquired (Fletcher, 1975), and, thus, suggest how content should be presented. But for more complex learning adequate operational statements are lacking. The majority of published work related to the learning process, while intellectually rewarding, is of little assistance in defining precise models. But, progress on this front is being made by Norman and associates at the Center for Human Information Processing at the University of California.

Norman and associates are engaged in modeling human memory. Their work has recently been described in the collection of papers: Explorations into Cognition (Norman & Rumelhart, 1975). At the foundation of the memory model employed by Norman et al. is a set of general functions which specify the relations which exist among subsets of concepts. The scheme for functional representation of related concepts borrows heavily from Fillmore's (1968) notion of "case grammar" in that it assigns particular roles to the concepts of a proposition.

The investigators are attempting to determine a primitive set of general functions so that all propositions can be expressed in terms of some subset of these more basic relations. So far their work has concentrated on describing functions that correspond to the meanings of verbs and the primitives which underlie verbal structures.

To illustrate the type of model Norman and coworkers are exploring we will consider the verb "give". The meaning of "give" that corresponds to one person placing an object in the possession of another person is represented by the following function:

give [agent, object, recipient, time]

The arguments of the function suggest the set of restrictions that apply to the values that can be used to fill the arguments. Thus the first argument--agent--denotes an animate being capable of instigating the action. This meaning of "give" can be further analyzed by considering the primitive notions which underlie the action, such as the notion of one event causing another and the notion of an object being transferred from some source to some goal.

The memory system Norman and coworkers have proposed adapts readily to a variety of cognitive tasks, including information acquisition (Rumelhart & Norman, 1975), language comprehension (Rumelhart & Levin, 1975) and problem solving (Eisenstadt & Kareev, 1975). It is not yet known whether the structures and processes modeled parallel human structures and processes, but the success of the model should encourage research in this area. Norman and associates have already pursued some research along this line (e.g., Gentner, 1975).

Modeling Teaching Strategies

CAI tutors in which the coding of procedures which control the flow of the instructional interaction is distinct from other modules of the tutor, such as the student or subject matter models, provide convenient tools for exploring various teaching scenarios. Typically the outward behavior of these tutors can be completely redesigned in just a matter of days. Hence the instructional researcher has at his disposal a well-controlled laboratory within which to model instructional strategies, implement the models, and compare their results in terms of learning achievement.

Separate control procedures exist in varying degrees in the systems being explored at BBN by Brown and Collins and in Atkinson's tutor of BASIC, BIP. Investigations concerned with isolating optimal instructional strategies are being pursued at both sites.

Collins and associates have engaged in modeling teaching strategies for their tutor of geography, GEO-SCHOLAR (Collins, 1976; Collins, Warnock & Passafiume, 1975). To isolate teaching strategies employed by human tutors, Collins and coworkers analyzed human tutor-tutee teaching sessions. Based on these analyses Collins and associates have proposed several hypotheses about how the tutor relates his teaching to the individual student. The

investigators have implemented some of these hypotheses as strategies which GEO-SCHOLAR can incorporate into its tutorial interactions. Experiments have been conducted to define those characteristics of implemented strategies which facilitate learning relative to a frame-oriented tutorial (Collins, 1974),

Research suggested by Brown and coworkers (e.g., Brown, 1975b; Brown & Burton, 1975) using tutors of problem solving like SOPHIE and the remedial algebra tutor promises to provide data concerning optimal strategies for teaching problem-solving techniques. Because of the computer's ability to search a solution space much more efficiently and effectively than can a human tutor, some strategies which can be studied easily in these problem-solving laboratories would be difficult, if not impossible, to model and evaluate in a noncomputerized setting.

Atkinson and his coworkers at the Institute for Mathematical Studies in the Social Sciences are exploring the use of curriculum networks to provide data for strategies which sequence instruction. By investigating problem sequencing as a function of student progress in skill acquisition, an attempt is being made to model optimal instructional sequencing strategies over the network (Barr, Beard, & Atkinson, 1975b).

CONCLUSIONS

The military must train a large, heterogeneous student population. Because the population is so varied, adaptive instructional environments are required for many training objectives. Because of the large number of students involved, such environments cannot be provided using labor intensive traditional approaches. By providing knowledge concerning the development of responsive automated tutors, ARPA supported projects have addressed these issues.

One large training area where cost-effective applications of responsive tutors seems most promising is tutorial simulations of laboratory situations. The traditional laboratory has severe limitations. Instructors may not have the time to adequately pursue a student's hypotheses or questions, or evaluate his procedures; in laboratories requiring measurements, a student's time may be ineffectually lost in preparations which add little to the training experience. If expensive equipment is being used or expended, as in combat training, the costs of training may severely constrain its effectiveness. Responsive CAI promises to eliminate these threats to learning effectiveness while possibly decreasing training costs by providing an environment

in which the student may manipulate, with "expert" supervision, a computer model of the system under investigation.

As test beds for instructional research, the systems developed by ARPA supported programs will aid in isolating software capabilities which facilitate training objectives. Work must then turn toward generalizing these capabilities to insure their transportability to other CAI environments. Also, as vehicles for instructional research, the systems will aid in identifying and rigorously defining optimal instructional strategies, thus providing guidelines for authors of CAI materials.

TRENDS

Advances in natural language comprehension will facilitate both the student-system and author-system interfaces, and, in turn increase the range of training objectives which can be achieved in a CAI environment.

As a CAI system becomes less responsive to natural language input, the number of extraneous skills a student must develop to interact with the program increases. The syntax of acceptable input statements becomes more prescribed, the semantics more constrained. In the framework of CAI, research in natural language processing will reduce interaction difficulties by decreasing the extent to which a student must process the system's language while increasing the extent to which the system processes the student's. Also because a broader range of student responses will be "understood" by the computer, the extent to which the system author must anticipate and constrain student inputs will decrease, encouraging the development of more sophisticated training environments.

By providing the system author with the opportunity to interact with the computer in natural language, the time needed to create CAI materials and the system limits imposed on those materials by programming language constraints will decrease. Presently much time is invested in learning programming skills and format conventions. The programming of complex CAI materials requires considerable programming experience and, hence, the availability of expert programmers. By providing the CAI author with a more natural interface, time needed to learn computer conventions and dependence on expert programming skills will decrease.

The capabilities being developed in responsive CAI tutors will be of value as job-performance aids. For example, BIP's student-program interpreter contains debugging aids that would be of assistance to programmers outside of the classroom environment (Barr & Beard, 1976). Also the simulation models developed for responsive tutors may be useful to authors of more traditional drill-and-practice or tutorial programs. Such models can quickly provide the author with data which would be time-consuming to collect by other means.

References

Anderson, J.R. and Bower, C.H. Human associative memory, New York: John Wiley and Sons, 1973.

Barr, A. and Beard, M. An instructional interpreter for BASIC. Proceedings for the symposium on computer science and education, 1976, 325-335.

Barr, A., Beard, M., and Atkinson, R.C. A rationale and description of a CAI program to teach the BASIC programming language. Instructional Science, 1975(a), 4, 1-31.

Barr, A., Beard, M., and Atkinson, R.C. The computer as a tutorial laboratory: The Stanford BIP project. (Technical Report No. 260). Stanford, Calif.: Institute for Mathematical Studies in the Social Sciences, Stanford University, 1975(b).

Beard, M., Barr, A., Fletcher, D. and Atkinson, R.C. The improvement and individualization of computer-assisted instruction: Final report. Stanford, Calif.: Institute for Mathematical Studies in the Social Sciences, Stanford University, 1975.

Brown, J.S. Progress report: Algebra tutoring system. Cambridge, Mass.: Bolt, Beranek and Newman, Inc., 1975(a).

Brown, J.S. Uses of artificial intelligence and advanced computer technology in education. Cambridge, Mass.: Bolt, Beranek and Newman, Inc., 1975(b).

Brown, J.S., Bell, A.G. and Zdybel, F. Steps toward an artifically intelligent CAI system: An interim report on research in progress. (Report No. UCI ICS 35). Irvine, Calif.: University of California, Irvine, 1973.

Brown, J.S. and Burton, R.R. Multiple representations of knowledge for tutorial reasoning. In D. Bobrow and A. Collins (Eds.) Representation and understanding: Studies in cognitive science, New York: Academic Press, Inc., 1975.

Brown, J.S., Burton, R.R. and Bell, A. Sophisticated instructional environment for teaching electronic troubleshooting. (Report No. AFHRL-TR-74-77). Irvine, Calif.: University of California, Irvine, 1974.

Brown, J.S., Burton, R.R., Bell, A.G. and Bobrow, R.J. SOPHIE: A sophisticated instructional environment. (Report No. AFHRL-TR-74-93). Cambridge, Mass.: Bolt, Beranek and Newman, Inc., 1974.

Bunderson, C.V. and Faust, G.W. Programmed and computer-assisted instruction. In N.L. Gage (Ed.). The psychology of teaching methods: The seventy-fifth yearbook of the National Society for the Study of Education. Chicago: The University of Chicago Press, 1976.

Burton, R.R. A semantically centered parsing system for mixed-initiative CAI systems. Paper presented at the Association for Computational Linguistics Conference, Amherst, Mass., July 1974.

Burton, R.R. and Brown, J.S. A tutoring and student modelling paradigm for gaming environments. Proceedings for the symposium on computer science and education, 1976, 236-247.

Carbonell, J.R. AI in CAI: An artificial-intelligence approach to computer-assisted instruction. IEEE Trans. on Man-Machine Systems, 1970(a), Vol. MMS-11, 190-202.

Carbonell, J.R. Mixed-initiative man-computer instructional dialogues. Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, 1970(b).

Carbonell, J.R. and Collins, A.M. Natural semantics in artificial intelligence. Proceedings of the third international joint conference on artificial intelligence, 1973, 344-351.

Collins, A.M. Comparison of two teaching strategies in computer-assisted instruction. (Report No. 2885). Cambridge, Mass.: Bolt, Beranek and Newman, Inc., 1974.

Collins, A.M. Processes in acquiring knowledge. In R.C. Anderson, R.J. Spiro, and W.E. Montague (Eds.). Schooling and the acquisition of knowledge. Hillsdale, N.J.: Erlbaum Assoc., 1976.

Collins, A.M. and Warnock, E.H. Semantic networks. (Report No. 2833). Cambridge, Mass.: Bolt, Beranek and Newman, Inc., 1974.

Collins, A.M., Warnock, E.H., Aiello, N. and Miller, M.L. Reasoning from incomplete knolwedge. In D. Bobrow and A. Collins (Ed.). Representation and understanding: Studies in cognitive science, New York: Academic Press, Inc., 1975.

Collins, A.M., Warnock, E.H. and Passafiume, J.J. Analysis and synthesis of tutorial dialogues. In G.H. Bower, The psychology of learning and motivation (Vol. 9). New York: Academic Press, Inc., 1975.

Eisenstadt, M. and Kareev, Y.A. Aspects of human problem solving: The use of internal representations. In D.A. Norman and D.E. Rumelhart (Eds.). Explorations in cognition. San Francisco: W.H. Freeman and Company, 1975.

Feigenbaum, E.A. and Feldman, J. (Eds.). Computers and thought, New York: McGraw-Hill Book Company, 1963.

Fillmore, C. The case for case. In Bach and Harms (Eds.). Universals in linguistic theory, Chicago, Ill.: Holt, 1968.

Fletcher, J.D. Models of the learner in computer-assisted instruction. (Report No. NPRDC TR 76-23). San Diego, Calif.: Navy Personnel Research and Development Center, 1975.

Gentner, D. Evidence for the psychological reality of semantic components: The verbs of possession. In D.A. Norman and D.E. Rumelhart (Eds.). Explorations in cognition. San Francisco: W.H. Freeman and Company, 1975.

Goldberg, A. Computer-assisted instruction: The application of theorem-proving to adaptive response analysis. (Technical Report No. 203). Stanford, Calif.: Institute for Mathematical Studies in the Social Sciences, Stanford University, 1973.

Grignetti, M.C., Gould, L., Hausmann, C.L., Bell, A.G., Harris, G. and Passafiume, J. Mixed-initiative tutorial system to aid users of the on-line system (NLS). (Report No. ESD-TR-75-58). Cambridge, Mass.: Bolt, Beranek and Newman, Inc., 1974.

Grignetti, M.C. and Warnock, E.H. Mixed-initiative information system for computer-aided training and decision-making. (Report No. ESD-TR-73-290). Cambridge, Mass.: Bolt, Beranek and Newman, Inc., 1973.

Kimball, R.B. Self-optimizing computer-assisted tutoring: Theory and practice. (Technical Report No. 206). Stanford, Calif.: Institute for Mathematical Studies in the Social Sciences, Stanford University, 1973.

Kingsley, E.H. and Stelzer, J. A theoretical basis for individualized instruction. (Report No. AFHRL-TR-74-10). Alexandria, Virginia: Human Resources Research Organization, 1974.

Koffman, E.B. and Blount, S.E. A modular system for generative CAI in machine-language programming. IEEE Trans. on Systems, Man, and Cybernetics, 1974, Vol. SMC-4, 335-343.

Koffman, E.B. and Blount, S.E. Artificial intelligence and automatic programming in CAI. Artificial Intelligence, 1975, 6, 215-234.

Nilsson, N.J. Artificial intelligence. Paper presented at the IFIP Congress 74, Stockholm, Sweden, August 1974.

Norman, D.A. Memory, knowledge, and the answering of questions. In R.L. Solso (Ed.). Contemporary issues in cognitive psychology: The Loyola symposium. Washington, D.C.: V.H. Winston and Sons, 1973.

Norman, D.A. Studies of learning and self-contained educational systems, 1973-1976. (Report No. 7601). LaJolla, Calif.: Center for Human Information Processing, University of California, San Diego, 1976.

Norman, D.A. and Rumelhart, D.E. Explorations in cognition, San Francisco: W.H. Freeman and Company, 1975.

Quillian, M.R. Semantic memory. In M. Minsky (Ed.). Semantic information Processing, Cambridge, Mass.: MIT Press, 1968.

Rockart, J.F. and Scott Morton, M. Computers and the learning process in higher education, New York: McGraw-Hill Book Company, 1975.

Rumelhart, D.E. Notes on a schema for stories. In D. Bobrow and A. Collins (Eds.). Representation and understanding: Studies in cognitive science, New York: Academic Press, Inc., 1975.

Rumelhart, D.E. Understanding and summarizing brief stories. (Report No. 58). LaJolla, Calif.: Center for Human Information Processing, University of California, San Diego, 1976.

Rumelhart, D.E. and Norman, D.A. The computer implementation. In D.A. Norman and D.E. Rumelhart (Eds.). Explorations in cognition. San Francisco: W.H. Freeman and Company, 1975.

Rumelhart, D.E. and Levin, J.A. A language comprehension system. In D.A. Norman and D.E. Rumelhart (Eds.). Explorations in cognition. San Francisco: W.H. Freeman and Company, 1975.

Self, J.A. Student models in computer-aid instruction. International Journal of Man-Machine Studies, 1974, 6, 261-276.

Siklossy, L. Computer tutors that know what they teach. Proceedings of the fall joint computer conference, 1970, 251-255.

Simmons, R.F. Semantic networks: Their conceptualization and use for understanding English sentences. In R.C. Schank and K.M. Colby (Eds.). Computer models of thought and language. San Francisco: W.H. Freeman and Company, 1973.

Uttal, W., Pasich, T., Rogers, M. and Hieronymus, R. Generative computer assisted instruction. Mental Health Research Institute, University of Michigan, Ann Arbor, Commun. 243, January 1969.

Wexler, J.D. Information networks in generative computer-assisted instruction. IEEE Trans. on Man-Machine Systems, 1970, Vol. MMS-11, 181-190.

Woods, W. Transition network grammars for natural language analysis. Communications of the ACM, 1970, 13, 591-602.

Woods, W.A. What's in a link: Foundations for semantic networks. In D. Bobrow and A. Collins (Eds.). Representation and understanding: Studies in cognitive science, New York: Academic Press, Inc., 1975.